(REV 10/01) U S DEPARTMENT OF COMMERCE PATENT AND TRADEMARK OFFICE

TRANSMITTAL LETTER TO THE UNITED STATES DESIGNATED/ELECTED OFFICE (DO/EO/US) CONCERNING A FILING UNDER 36 U.S.C. 371

ATTORNEY'S DOCKET NUMBER: COLB-121XX

U S APPLICATION NO. (If known, see 37 CFR 1.5): 10/070594

INTERNATIONAL APPLICATION NO.	INTERNATIONAL FILING DATE	PRIORITY DATE CLAIMED
PCT/IL00/00540	07 September 2000 (07.09.2000)	08 September 1999 (08.09.2000)

TITLE OF INVENTION: SYNCHRONIZATION OF INTERRUPTS WITH DATA PACKETS

APPLICANT(S) FOR DO/EO/US: Michael Kagan, Diego Crupnicoff, Freddy Gabbay, Shimon Rottenberg

Applicant herewith submits to the United States Designated/Elected Office (DO/EO/US) the following items and other information:

1. ☒ This is a **FIRST** submission of items concerning a filing under 35 U.S.C. 371.
2. ☐ This is a **SECOND** or **SUBSEQUENT** submission of items concerning a filing under 35 U.S.C. 371.
3. ☐ This express request to begin national examination procedures (35 U.S.C. 371(f)). The submission must include items (5), (6), (9) and (21) indicated below.
4. ☒ The US has been elected by the expiration of 19 months from the priority date (Article 31).
5. ☒ A copy of the International Application as filed (35 U.S.C. 371(c)(2)) **Published Appln. WO 01/18654 (in English)**
   a. ☒ is attached hereto (required only if not transmitted by the International Bureau).
   b. ☐ has been communicated by the International Bureau.
   c. ☐ is not required, as the application was filed in the United States Receiving Office (RO/US).
6. ☐ An English language translation of the International Application as filed (35 U.S.C. 371(c)(3)).
   a. ☐ is attached hereto.
   b. ☐ had been previously submitted under 35 U.S.C. 154(d)(4).
7. ☒ Amendments to the claims of the International Application under PCT Article 19 (35 U.S.C. 371(c)(3))
   a. ☐ are attached herewith (required only if not transmitted by the International Bureau).
   b. ☐ have been communicated by the International Bureau.
   c. ☐ have not been made; however, the time limit for making such amendments has NOT expired.
   d. ☒ have not been made and will not be made.
8. ☐ An English language translation of the amendments to the claims under PCT Article 19 (35 U.S.C. 371(c)(3)).
9. ☐ An oath or declaration of the inventor(s) (35 U.S.C. 371(c)(4)).
10. ☐ A translation of the annexes to the International Preliminary Examination Report under PCT Article 36 (35 U.S.C. 371(c)(5)).

Items 11. to 20. below concern document(s) or information included:

11. ☒ An Information Disclosure Statement under 37 CFR 1.97 and 1.98. (with attached copy of International Search Report)
12. ☐ An assignment document for recording. A separate cover sheet in compliance with 37 CFR 3.28 and 3.31 is included.
13. ☒ A FIRST preliminary amendment.
14. ☐ A SECOND or SUBSEQUENT preliminary amendment.
15. ☐ A substitute specification.
16. ☐ A change of power of attorney and/or address letter.
17. ☐ A computer readable form of the sequence listing in accordance with PCT Rule 13ter.2 and 35 U.S.C. 1.821-1.825
18. ☐ A second copy of the published international application under 35 U.S.C. 154(d)(4).
19. ☐ A second copy of the English language translation of the international application under 35 U.S.C. 154(d)(4).
20. ☒ Other items or information:
    FORMAL DRAWINGS (2 sheets)

Express Mail Number
EV 009953392 US

U.S. APPLICATION NO. (If known, see 37 CFR 1.5)	INTERNATIONAL APPLICATION NO	ATTORNEY'S DOCKET NUMBER
10/070594	PCT/IL00/00540	COLB-121XX

21. ☒ The following fees are submitted:

BASIC NATIONAL FEE (37 CFR 1.492 (a) (1) - (5)):		**CALCULATIONS** PTO USE ONLY	
Neither international preliminary examination fee (37 CFR 1.482) nor international search fee (37 CFR 1.445(a)(2)) paid to USPTO and International Search Report not prepared by the EPO or JPO	**$1,040.00**		
International preliminary examination fee (37 CFR 1.482) not paid to USPTO but International Search Report prepared by the EPO or JPO	**$890.00**		
International preliminary examination fee (37 CFR 1.482) not paid to USPTO but international search fee (37 CFR 1.445(a)(2)) paid to USPTO	**$740.00**		
International preliminary examination fee paid to USPTO (37 CFR 1.482) but all claims did not satisfy provisions of PCT Article 33(1)-(4)	**$710.00**		
International preliminary examination fee paid to USPTO (37 CFR 1.482) and all claims satisfied provisions of PCT Article 33(1)-(4)	**$100.00**		
ENTER APPROPRIATE BASIC FEE AMOUNT =		$ 710.00	
Surcharge of **$130.00** for furnishing the oath or declaration later than ☐ 20 ☐ 30 months from the earliest claimed priority date (37 CFR 1.492(e)).		$ 0	

CLAIMS	NUMBER FILED	NUMBER EXTRA	RATE		
Total claims	27 - 20 =	7	X **$18.00**	$ 126.00	
Independent claims	3 - 3 =	0	X **$84.00**	$ 0	
MULTIPLE DEPENDENT CLAIM(S) (if applicable)			+**$280.00**	$ 0	
TOTAL OF ABOVE CALCULATIONS =				$ 836.00	
☒ Applicant claims small entity status. See 37 CFR 1.27. The fees indicated above are reduced by 1/2.				$ 418.00	
SUBTOTAL =				$ 418.00	
Processing fee of **$130.00** for furnishing the English translation later than ☐ 20 ☐ 30 months from the earliest claimed priority date (37 CFR 1.492(f)).				$ 0	
TOTAL NATIONAL FEE =				**$ 418.00**	
Fee for recording the enclosed assignment (37 CFR 1.21(h)). The assignment must be accompanied by an appropriate cover sheet (37 CFR 3.28, 3.31). **$40.00** per property				$	
TOTAL FEES ENCLOSED =				$ 418.00	
				Amount to be Refunded:	$
				Charged:	$

a. ☒ A check in the amount of $ 418.00 to cover the above fees is enclosed. A check in the amount of $_____ is enclosed for the assignment recordation fee.

b. ☐ Please charge my Deposit Account No. _____ in the amount of $_____ to cover the above fees. A duplicate copy of this sheet is enclosed.

c. ☒ The Commissioner is hereby authorized to charge any additional fees which may be required, or credit any overpayment to Deposit Account No. 23-0804 . A duplicate copy of this sheet is enclosed.

NOTE: Where an appropriate time limit under 37 CFR 1.494 or 1.495 has not been met, a petition to revive (37 CFR 1.137(a) or (b)) must be filed and granted to restore the application to pending status.

Customer Number 207
SEND ALL CORRESPONDENCE TO:

Weingarten, Schurgin, Gagnebin & Lebovici LLP
Ten Post Office Square
Boston, Massachusetts 02109

SIGNATURE

NAME: Charles L. Gagnebin III
REGISTRATION NUMBER: 25,467

Date: 3-8-2

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application	:	MICHAEL KAGAN, ET AL.
Application No.	:	
Filed	:	Herewith
For	:	SYNCHRONIZATION OF INTERRUPTS WITH DATA PACKETS
Examiner	:	
Attorney's Docket	:	COLB-121XX

Group Art Unit:

* *

I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail in an envelope addressed to: Commissioner for Patents, Washington, D.C. 20231 on _____.

By:_____
Charles L. Gagnebin III
Registration No. 25,467
Attorney for Applicant(s)

* *

PRELIMINARY AMENDMENT

BOX PCT
Commissioner for Patents
Washington, D.C. 20231

Sir:

Kindly enter the following Preliminary Amendment in the above-identified application:

In the Claims:

PLEASE CANCEL CLAIMS 1-27.

WEINGARTEN, SCHURGIN,
GAGNEBIN & LEBOVICI LLP
TEL. (617) 542-2290
FAX. (617) 451-0313

Express Mail Number
EV 009953392 US

Please add the following new claims 28-54:

28. A method for communication between a peripheral device and a central processing unit (CPU), comprising:

receiving data from the peripheral device for transmission to a memory associated with the CPU;

receiving an interrupt signal from the peripheral device associated with the data;

sending one or more data packets containing the data over a switched serial connection to a host interface serving the memory and the CPU; and

sending an interrupt packet over the switched serial connection to the host interface, responsive to which an interrupt input of the CPU is asserted only after the one or more data packets have arrived at the host network interface.

29. A method according to claim 28, wherein receiving the data comprises receiving parallel data over a local bus from the peripheral device.

30. A method according to claim 28, wherein receiving the data comprises receiving data to be written to the memory by direct memory access.

WEINGARTEN, SCHURGIN,
GAGNEBIN & LEBOVICI LLP
TEL. (617) 542-2290
FAX. (617) 451-0313

31. A method according to claim 28, wherein sending the interrupt packet comprises reading a cause of the interrupt from the peripheral device, and incorporating the cause in the interrupt packet.

32. A method according to claim 31, and comprising receiving the interrupt packet at the host interface, and writing the cause to a predetermined address in the memory, to be read by the CPU after the interrupt input is asserted.

33. A method according to claim 28, wherein sending the interrupt packet comprises sending the interrupt packet after receiving an acknowledgment from the memory that the data have been written thereto.

34. A method according to claim 28, wherein sending the one or more data packets comprises sending the data packets over a selected lane through a packet-switched network, and wherein sending the interrupt packet comprises sending the interrupt packet over the selected lane following the data packets.

35. A method according to claim 28, and comprising:

WEINGARTEN, SCHURGIN,
GAGNEBIN & LEBOVICI LLP
TEL. (617) 542-2290
FAX. (617) 451-0313

receiving the data packets and the interrupt packet at the host interface;

conveying the data in the packets for delivery to the memory over a local bus coupling the host interface to the memory and the CPU; and

notifying the CPU when all of the data have been conveyed.

36. A method according to claim 35, wherein conveying the data in the packets comprises passing the data to a system controller on the bus, and wherein notifying the CPU comprises informing the CPU when an acknowledgment is received by the host interface from the system controller.

37. A method according to claim 36, wherein informing the CPU comprises asserting the interrupt input of the CPU after the acknowledgment from the system controller has been received.

38. A method according to claim 35, wherein notifying the CPU comprises asserting the interrupt input of the CPU responsive to receiving the interrupt packet at the host network interface.

39. Communication apparatus, comprising:

WEINGARTEN, SCHURGIN,
GAGNEBIN & LEBOVICI LLP
TEL. (617) 542-2290
FAX. (617) 451-0313

a serial interface, which is operative to receive data from a peripheral device and to transmit the data in the form of one or more data packets via a switched serial connection to a host interface, for writing to a memory associated with a central processing unit (CPU) served by the host interface; and

an interrupt processor, adapted to send an interrupt packet over the switched serial connection to the host interface to signal that the data have been transmitted, thus causing an interrupt input of the CPU to be asserted only after the one or more data packets have arrived at the host network interface.

40. Apparatus according to claim 39, wherein the serial interface is adapted to receive the data from the peripheral device over a local parallel bus.

41. Apparatus according to claim 39, wherein the interrupt processor is operative to receive a cause of the interrupt from the peripheral device, and to incorporate the cause in the interrupt packet.

42. Apparatus according to claim 41, wherein the host interface is adapted to receive the interrupt packet and to write the

WEINGARTEN, SCHURGIN,
GAGNEBIN & LEBOVICI LLP
TEL. (617) 542-2290
FAX. (617) 451-0313

cause to a predetermined address in the memory, to be read by the CPU after the interrupt input is asserted.

43. Apparatus according to claim 39, wherein the interrupt processor is adapted to send the interrupt packet after receiving an acknowledgment from the memory that the data have been written thereto.

44. Apparatus according to claim 39, wherein the serial interface is coupled to send the data packets over a selected lane through the network, and wherein the processor is adapted to send the interrupt packet over the selected lane following the data packets.

45. Apparatus according to claim 44, wherein the switched serial connection comprises a switch having a receive queue into which the serial interface places the data packets, and wherein the interrupt processor is adapted to place the interrupt packet into the receive queue following the data packets.

46. Apparatus according to claim 39, wherein the host interface is coupled to receive the data and interrupt packets transmitted over the switched serial connection, and is operative to convey

WEINGARTEN, SCHURGIN,
GAGNEBIN & LEBOVICI LLP
TEL. (617) 542-2290
FAX. (617) 451-0313

the data in the packets for delivery to the memory over a local bus coupled to the memory and the CPU and to notify the CPU when all of the data have been conveyed.

47. Apparatus according to claim 46, wherein the host interface is coupled to assert the interrupt to the CPU responsive to the interrupt packet.

48. Apparatus according to claim 39, wherein the switched serial connection is part of a switching fabric.

49. Communication apparatus, comprising:

a host adapter, which is operative to receive data packets transmitted over a switched serial connection from a peripheral device, and to convey data from the packets for delivery to a memory associated with a CPU over a local bus that is coupled to the memory and the CPU, and further to receive an interrupt packet sent over the switched serial connection responsive to an interrupt signal asserted by the peripheral device after sending the data to the network; and

a host interface processor, adapted, responsive to the interrupt packet, to notify the CPU when all of the data have been conveyed to the local bus.

WEINGARTEN, SCHURGIN,
GAGNEBIN & LEBOVICI LLP
TEL. (617) 542-2290
FAX. (617) 451-0313

50. Apparatus according to claim 49, wherein the host adapter is operative to convey the data to the memory by direct memory access.

51. Apparatus according to claim 49, wherein the host adapter is operative to convey the data to a system controller on the bus, and wherein the CPU is notified when an acknowledgment is received by the host adapter from the system controller.

52. Apparatus according to claim 51, wherein the host interface processor is coupled to assert the interrupt input of the CPU after the acknowledgment from the system controller has been received.

53. Apparatus according to claim 49, wherein the host interface processor is coupled to assert the interrupt input of the CPU responsive to receipt of the interrupt packet by the host adapter.

54. Apparatus according to claim 49, wherein the switched serial connection is part of a switching fabric.

WEINGARTEN, SCHURGIN,
GAGNEBIN & LEBOVICI LLP
TEL. (617) 542-2290
FAX. (617) 451-0313

Attorney Docket No. COLB-121XX
Filed: Herewith
Group Art Unit:

REMARKS

This application contains new claims 28-54. No new matter has been introduced. Favorable consideration of the new claims is respectfully requested.

Claims 28-54 hereby added to the present patent application are substantially parallel and similar in substance to original claims 1-27, which are now cancelled. A number of changes have been introduced in the claims, however, in order to more clearly recite aspects of the present invention. Specifically, references to a "packet-switching network" in the original claims have been replaced by references to switched serial connections. Connections of this sort are characteristic of switch fabrics, *inter alia*, such as the InfiniBand fabric described in the specification of the present patent application. In the same vein, for the sake of clarity, the term "target channel adapter" used in claim 12 has been replaced by a serial interface in claim 39; and "host channel adapter" in claim 22 has been replaced by a host adapter in claim 49. With regard to claims 21 and 27, the proprietary name "InfiniBand" has been replaced with its generic alternative, a switching fabric, in new claims 48 and 54.

WEINGARTEN, SCHURGIN,
GAGNEBIN & LEBOVICI LLP
TEL. (617) 542-2290
FAX. (617) 451-0313

In the International Preliminary Examination Report (IPER), the Examiner indicated that claims 1-3, 7, 8, 11-13, 17, 19, 20 and 22-26 lack novelty over Gentry et al. (US 5,659,758). Applicant respectfully traverses this rejection. Gentry describes an interrupt modulator, which generates interrupts upon receiving packets over a communication network. This interrupt modulator passes the interrupts directly to the processor that is to receive the packet data, as shown in Gentry's Fig. 3. Although the packets themselves are received from a high-speed communication network by a packet extractor, Gentry neither teaches nor suggests that the interrupts may be sent in packet form over the network, as provided by preferred embodiments of the present invention. Therefore, applicant respectfully submits that a method of communication based on sending an interrupt packet over a switched serial connection to the host interface, as recited in new claim 28, is patentable over Gentry. New independent claims 39 and 49 contain similar recitations. In view of the patentability of these independent claims, the new dependent claims in the present application are believed to be patentable, as well.

In view of the above amendments and remarks, applicant respectfully submits that all of the claims in the present

WEINGARTEN, SCHURGIN,
GAGNEBIN & LEBOVICI LLP
TEL. (617) 542-2290
FAX. (617) 451-0313

Attorney Docket No. COLB-121XX
Filed: Herewith
Group Art Unit:

application are in condition for allowance. Prompt notice to this effect is respectfully requested.

The Examiner is encouraged to telephone the undersigned attorney to discuss any matter which would expedite allowance of the present application.

Respectfully submitted,

MICHAEL KAGAN, ET AL.

By: _____
Charles L. Gagnebin III
Registration No. 25,467
Attorney for Applicant(s)

WEINGARTEN, SCHURGIN, GAGNEBIN
& LEBOVICI LLP
Ten Post Office Square
Boston, MA 02109
Telephone: (617) 542-2290
Telecopier: (617) 451-0313

Date: 3-8-2

CLG/mc/268698-1

WEINGARTEN, SCHURGIN,
GAGNEBIN & LEBOVICI LLP
TEL. (617) 542-2290
FAX. (617) 451-0313

SYNCHRONIZATION OF INTERRUPTS WITH DATA PACKETS

CROSS-REFERENCE TO RELATED APPLICATION

This application claims the benefit of U.S. Provisional Patent Application 60/152,849, filed September 8, 1999, and of U.S. Provisional Patent Application 60/175,339, filed January 10, 2000. Both of these co-pending applications are assigned to the assignee of the present patent application and are incorporated herein by reference.

FIELD OF THE INVENTION

The present invention relates generally to computing systems, and specifically to systems that use packet-switching fabrics to connect a computer host to peripheral devices.

BACKGROUND OF THE INVENTION

In current-generation computers, the central processing unit (CPU) is connected to the system memory and to peripheral devices by a parallel bus, such as the ubiquitous Peripheral Component Interface (PCI) bus. As data path-widths grow, and clock speeds become faster, however, the parallel bus is becoming too costly and complex to keep up with system demands. In response, the computer industry is moving toward fast, packetized, serial input/output (I/O) bus architectures, in which computing hosts and peripheral are linked by a switching network, commonly referred to as a switching fabric. A number of architectures of this type have been proposed, including "Next Generation I/O" (NGIO) and "Future I/O" (FIO), culminating in the "InfiniBand" architecture, which has been advanced by a consortium led by a group of industry leaders (including Intel, Sun, Hewlett Packard, IBM, Compaq, Dell and Microsoft). Storage Area Networks (SAN) provide a similar, packetized, serial approach to high-speed storage access, which can also be implemented using an InfiniBand fabric.

In a parallel bus-based computer system, when a peripheral device needs to deliver data to the CPU, it typically writes the data to the memory over the bus, using direct memory access. When the peripheral has finished writing, it asserts an interrupt to the CPU on one of the interrupt lines of the bus. Bus arbitration ensures that the CPU will not attempt to read the data from the memory until the writing of the data is complete. On the other hand, when the peripheral device and the CPU are connected by a packet-switching fabric, such as an InfiniBand fabric, they operate asynchronously. Furthermore, the data sent to the memory and the interrupt to the CPU travel over different paths, or channels. Typically, a separate line or

channel is provided to connect the interrupt pin of the peripheral device to an interrupt controller of the CPU, bypassing the switching fabric. Therefore, there is no *a priori* assurance that all of the data will have been written to the memory before the CPU begins reading.

The "race" between the interrupt path and the data path can result in errors (as when a CPU read stalls the data). Care must therefore be taken to synchronize data and interrupt handling and to make sure that the data have been completely written to the memory before the CPU attempts to read it.

A common solution in this situation is to program the CPU to access the peripheral device before accessing the memory, typically by performing a "configuration read" from the peripheral device. In this mode of operation, after the peripheral device has asserted the interrupt to the CPU (indicating that the last item of data has been sent to the memory), the CPU issues a read request through the switching fabric, to read an interrupt cause register in the peripheral device. The peripheral device responds to the read request by sending a packet containing the interrupt cause to the CPU over the same channel as it used to send the data to the memory. Since packets are ordered within a channel, the response to configuration read arrives at the CPU after all of the previous writes have been flushed to memory. The CPU begins to read the data from the memory only after it has received the interrupt cause packet back from the peripheral device. The configuration read thus serves two crucial purposes: it provides the CPU with the cause information that it needs in order to serve the interrupt, and it ensures that the CPU reads the memory only after all of the data have been written there.

This scheme has a number of serious performance drawbacks, however. Every interrupt sent by the peripheral device necessitates an additional exchange of messages through the switching fabric between the CPU and peripheral device. The exchange adds substantial latency – typically 10 microseconds or more - every time the CPU must service an interrupt. Furthermore, since configuration reads are used as synchronization barriers, the CPU is stalled from the moment the configuration read request is issued until its response has arrived. Valuable CPU time is therefore wasted waiting for the interrupt cause to be retrieved.

U.S. Patent 5,689,713, whose disclosure is incorporated herein by reference, describes a method for interrupt request handling in a packet-switched computer system. The system may include a number of interrupt sources, which direct interrupts to any of a number of interrupt handlers. A system controller acts as an intermediary between interrupting devices and

"interruptees." It includes an interrupt queue coupled to each interrupt source for receiving multiple interrupt requests, and an output queue coupled to each interrupt handler The controller thus enables asynchronous data from multiple sources to be conveyed across a packet-switched interconnection, while providing a dedicated channel for interrupts associated with the data packets.

SUMMARY OF THE INVENTION

It is an object of the present invention to provide an improved method and system for passing data packets and associated interrupts through a switching fabric.

It is a further object of some aspects of the present invention to provide a method and system for communication between a CPU and peripheral devices via a switching fabric that ensures proper synchronization between data and interrupts transmitted over the fabric.

It is still a further object of some aspects of the present invention to provide a method and system for communication between a CPU and peripheral devices via a switching fabric that reduces latency and processing time required for servicing of interrupts by the CPU.

In preferred embodiments of the present invention, a CPU and a peripheral device are linked to a packet- switching fabric by respective host and target network interfaces. The target interface receives data over a local bus from the peripheral device, for transmission in the form of packets to a system memory associated with the CPU. After sending the data, the peripheral device asserts an interrupt. The interrupt from the device is connected to an interrupt input of the target interface, rather than directly to the CPU or to a central system controller, as in systems known in the art. In response to the interrupt, the target interface reads the interrupt cause from the peripheral device, and then sends a special interrupt packet, including the interrupt cause, to the host interface. Preferably, the target interface sends the interrupt packet on the same channel as it sent the data packets, i.e., over the same "virtual lane," or route, and with the same priority as the data packets. It thus assures that the host interface will receive the interrupt packet only after it has received all of the preceding data packets.

Upon receiving the interrupt packet, the host interface places the interrupt cause in a predefined register in the memory. An interrupt signal is then sent from the host interface to an interrupt input of the CPU. Upon receiving the signal, the CPU checks to ensure that the host interface has finished writing all of the data from the peripheral device to the memory. This

check serves a similar purpose to the configuration read described in the Background of the Invention. Only after completing the check does the CPU read the interrupt cause and begin processing the data in the memory. The CPU performs all of these steps locally, communicating with the host interface and memory over a local system bus, with latency on the order of nanoseconds, rather than having to exchange messages with the peripheral device through the switching fabric, taking many microseconds. As a result, interrupt response latency is minimized, and the CPU does not waste precious time and resources waiting for the configuration read response.

In preferred embodiments of the present invention, the switching fabric comprises an InfiniBand network, and the host and target interfaces respectively comprise host and target channel adapters. It will be appreciated, however, that the principles of the present invention may similarly be applied to transmission of interrupts through substantially any packet-switched network.

There is therefore provided, in accordance with a preferred embodiment of the present invention, a method for conveying data over a packet-switching network, including:

receiving data from a peripheral device for transmission via the network to a memory associated with a central processing unit (CPU);

receiving an interrupt signal from the peripheral device associated with the data;

sending one or more data packets containing the data over the network to a host network interface serving the memory and the CPU; and

sending an interrupt packet over the network to the host network interface, responsive to which an interrupt input of the CPU is asserted only after the one or more data packets have arrived at the host network interface.

Typically, receiving the data includes receiving parallel data over a local bus from the peripheral device. Additionally or alternatively, receiving the data includes receiving data to be written to the memory by direct memory access.

Preferably, sending the interrupt packet includes reading a cause of the interrupt from the peripheral device, and incorporating the cause in the interrupt packet. Further preferably, the method includes receiving the interrupt packet at the host network interface, and writing the cause to a predetermined address in the memory, to be read by the CPU after the interrupt input is asserted.

In a preferred embodiment, sending the interrupt packet includes sending the interrupt packet after receiving an acknowledgment from the memory that the data have been written thereto.

Preferably, sending the one or more data packets includes sending the data packets over a selected channel through the network, and sending the interrupt packet includes sending the interrupt packet over the selected channel following the data packets.

Further preferably, the method includes:

receiving the data packets and the interrupt packet at the host network interface;

conveying the data in the packets for delivery to the memory over a local bus coupling the host network interface to the memory and the CPU; and

notifying the CPU when all of the data have been conveyed.

Most preferably, conveying the data in the packets includes passing the data to a system controller on the bus, and notifying the CPU includes informing the CPU when an acknowledgment is received by the host network interface from the system controller, typically by asserting the interrupt input of the CPU after the acknowledgment from the system controller has been received. Additionally or alternatively, notifying the CPU includes asserting the interrupt input of the CPU responsive to receiving the interrupt packet at the host network interface.

There is also provided, in accordance with a preferred embodiment of the present invention, network interface apparatus, including:

a target channel adapter, which is operative to receive data from a peripheral device for transmission via a packet-switching network to a memory associated with a central processing unit (CPU) and to send one or more data packets containing the data over the network to a host network interface serving the memory and the CPU; and

a target interface processor, adapted to receive an interrupt signal from the peripheral device associated with the data, and to send an interrupt packet over the network to the host network interface, responsive to which an interrupt input of the CPU is asserted only after the one or more data packets have arrived at the host network interface.

There is further provided, in accordance with a preferred embodiment of the present invention, network interface apparatus, including:

a host channel adapter, which is operative to receive data packets transmitted over a packet-switching network from a peripheral device, and to convey data from the packets for delivery to a memory associated with a CPU over a local bus that is coupled to the memory and the CPU, and further to receive an interrupt packet sent over the network responsive to an interrupt signal asserted by the peripheral device after sending the data to the network; and

a host interface processor, adapted, responsive to the interrupt packet, to notify the CPU when all of the data have been conveyed to the local bus.

Preferably, the target and host channel adapters include InfiniBand adapters.

The present invention will be more fully understood from the following detailed description of the preferred embodiments thereof, taken together with the drawings in which:

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram that schematically illustrates a computing system based on a packet- switching fabric, in accordance with a preferred embodiment of the present invention;

Fig. 2 is a flow chart that schematically illustrates a method for transmitting data from a peripheral device to a CPU in the system of Fig. 1, in accordance with a preferred embodiment of the present invention; and

Fig. 3 is a flow chart that schematically illustrates a method for processing data received by the CPU in the system of Fig. 1, in accordance with a preferred embodiment of the present invention.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

Fig. 1 is a block diagram that schematically illustrates a computing system 20 built around a switching fabric 26, in accordance with a preferred embodiment of the present invention. The switching fabric preferably comprises an InfiniBand fabric, as described in the Background of the Invention, and some of the terms used hereinbelow are specific to the InfiniBand architecture. It will be understood, however, that the system architecture and methods of communication described herein are in no way limited to InfiniBand, and that other switching fabrics, as are known in the art, may be configured to handle and convey interrupts in a similar manner.

A CPU 21 is coupled to communicate via a system bus 52 with a system controller 24 and a system memory 22, as is known in the art. Typically (although not necessarily), the CPU comprises an Intel Pentium processor, and bus 52 is a proprietary bus used in conjunction with

this processor. System controller 24 is coupled to a standard I/O bus 50, such as a PCI bus, for the purpose of communicating with peripheral devices, such as I/O adapters of various types. One such peripheral device 25 is shown in Fig. 1 by way of example, but in practical applications, system 20 typically comprises multiple peripheral devices and, possibly, multiple CPUs. Peripheral device 25 includes an interrupt output 48, which it asserts in order to gain the attention of the CPU. In systems known in the art, interrupt output 48 is connected directly to an interrupt controller 38, such as an Intel 8259 device, which actuates an appropriate interrupt input 27 of CPU 21 when the interrupt is asserted. In system 20, however, interrupt output 48 and input 27 are linked only through fabric 26, as described hereinbelow.

Bus 50 is coupled to fabric 26 by a host network interface unit 28. This unit comprises a host channel adapter (HCA) 32, which interfaces with bus 50 and converts data between packet and parallel forms. Alternatively, the HCA may be designed to interface with system bus 52. A switch 30 links the HCA to one or more core switches in the fabric. Ordinarily, data in packets received by switch 30 from fabric 26 are passed through HCA 32 to bus 50. An exception is made, however, for management packets, which are packets that carry a special header identifying themselves as such and including a local identifier (LID) address of either switch 30 or HCA 32. These packets contain control instructions for the switch or HCA. They are placed in a dedicated register of the switch or HCA, as appropriate, which then attempts to decode the instructions and carry them out. Typically, the processing capabilities of the switch and HCA are very limited, and they are assisted by a fabric service agent (FSA), as described below, in dealing with at least some of these management packets.

A host interface unit controller 36 acts as the FSA in interface unit 28. The controller preferably comprises a microprocessor with random access memory (RAM) for software code and data, communicates with HCA 32 and switch 30. Alternatively, the controller may comprise a hard-wired hardware element or digital signal processor. When HCA 32 or switch 30 receives a management packet that it cannot decode, it passes the packet to the controller. The controller decodes the packet, preferably based on suitable software stored in its code RAM. It then takes whatever action is called for by the packet, such as giving appropriate instructions to HCA 32 or switch 30. When the HCA receives an interrupt packet, as described below, the actions taken by controller 36 also include signaling interrupt controller 38 via an interrupt output of unit 28, so as to actuate interrupt input 27 of CPU 21.

Although for simplicity, only a single interrupt line from unit 28 to controller 38 is shown in Fig. 1, the unit preferably comprises multiple interrupt lines These lines can be actuated selectively by controller 36 so as to send multiple, different interrupts to CPU 21 depending on the content of interrupt packets received by the HCA. Alternatively or additionally, the different interrupt lines may be used to signal other host devices that are linked to bus 50.

Peripheral device 25 is coupled to fabric 26 by a target network interface unit 40, similar in structure to unit 28. A target channel adapter (TCA) 42 in unit 40 interfaces via an I/O bus 53 with device 25. Typically, although not necessarily, bus 53 comprises a PCI bus, like bus 50. A switch 44 links the TCA to the switching fabric. A target unit controller 46, similar to controller 36, acts as FSA to TCA 42 and switch 44 and also has a suitable input to receive signals from interrupt output 48 of device 25.

Fig. 2 is a flow chart that schematically illustrates a method by which target interface unit 40 processes and transmits data from peripheral device 25 to HCA 32 over fabric 26, in accordance with a preferred embodiment of the present invention. At a data writing step 60, device 25 writes data via bus 53 to TCA 42, to be conveyed by direct memory access to memory 22. The peripheral device assigns a priority to the data to be transmitted and informs the TCA of this priority. At a data sending step 62, the TCA packetizes the data and sends it over fabric 26 to the address of HCA 32, with the priority assigned by the peripheral device. A packet header instructs the HCA to write the data to memory 22. Preferably, the TCA negotiates with switch 44 and fabric 26 to assign a fixed route for all of the packets through the fabric. Such a route, together with the priority of the packets, is referred to herein as a channel. InfiniBand specifies that packets travelling over the same channel are always kept in their original order.

When device 25 has finished posting to TCA 42 all of the data that it has to send, it asserts interrupt output 48, at an interrupt assertion step 64. At the same time, the peripheral device places the cause for the interrupt (in this case, to instruct CPU 21 to read the data from memory 22) in an interrupt cause register 49. In systems known in the art, when the CPU receives the interrupt, it must communicate with the peripheral device in order to read this register. In system 20, however, the interrupt signal is received by controller 46, which instructs TCA 42 to read the interrupt cause from register 49, at a cause reading step 66.

Based on the interrupt cause information read by the TCA, controller 46 constructs an interrupt packet containing the interrupt cause information, at an interrupt packet sending step 68. The interrupt packet is a management packet addressed to the LID of HCA 32. It is preferably sent by controller 46 over the same channel, or virtual lane, as the data packets, after the last of the data packets has been sent. The interrupt packet also identifies the data with which the interrupt is associated. As a result, when the interrupt packet arrives at its destination, controller 36 will be able to generate an interrupt to CPU 21 that is associated with the appropriate memory write, as described below. Controller 46 assures than interrupt packet is sent to the fabric after all of the data packets have already been accepted for sending. It thus ensures that HCA 32 will receive the interrupt packet only after it has received all of the data packets.

As an alternative, controller 46 may delay sending the interrupt packet until TCA 42 receives an acknowledgment from memory 22 that it has received all of the data. This approach introduces additional delay before CPU 21 can receive and act upon the interrupt, but it obviates the need to ensure that the interrupt packet is routed over the same channel as the data packets. Such an approach may be called for in particular when switching fabric 26 comprises a network in which consistent routing and ordering are not necessarily maintained among successive packets. This approach can also be used when the interrupt path and data path are not the same, and fork at an earlier stage than in Fig. 1. Such path incongruity may occur, for example, when the device writing data to the memory is different from the device asserting the interrupt to the CPU. Sometimes it is also desirable to send interrupts on different (high-priority) routes, because data routes can be congested, causing interrupt messages to get stuck behind data.

Fig. 3 is a flow chart that schematically illustrates a method by which data and accompanying interrupt packets are received and processed by host interface unit 28 and CPU 21, in accordance with a preferred embodiment of the present invention. At a packet reception step 70, HCA 32 receives the data and interrupt packets sent from target interface unit 40. The HCA posts the data in the data packets via bus 50 to a buffer 58 of system controller 24. The system controller proceeds to write the data from its buffer to the appropriate addresses in memory 22, as is known in the art. The HCA passes the interrupt packet to controller 36 for

decoding, at an interrupt processing step 72. The controller extracts the cause of the interrupt and posts this information, via HCA 32, to an interrupt cause register 56 in memory 22.

Before CPU 21 services the interrupt represented by the interrupt packet, it is necessary to ensure that all of the associated data have been written to memory 22, at a delivery completion step 74. In the case that controller 46 of target interface unit 40 is programmed to send the interrupt packet only after receiving the acknowledgment from memory 22, as described above, this problem is already solved. Otherwise, controller 36 preferably waits to assert the interrupt until system controller 24 has acknowledged to HCA 32 that it has received all of the data. In response to this acknowledgment, controller 36 sends an interrupt signal to interrupt controller 38, at an interrupt assertion step 76. The interrupt controller actuates interrupt input 27 of CPU 21, to inform the CPU that an interrupt has arrived from HCA 32. In response to the interrupt, the CPU preferably sends a dummy read command to the HCA, in order to ensure that buffer 58 is flushed to memory 22 before the CPU itself begins to process the data in the memory.

As a further alternative, as long as it is assured that the interrupt packet reached HCA 32 after the last of the data packets (which will be the case when all of the packets are sent over the same channel, as described above), controller 36 may send the interrupt signal to interrupt controller 38 immediately, without waiting for an acknowledgment from system controller 24. In this case, upon receiving the interrupt, CPU 21 preferably sends a "fence" command to HCA 32. This command instructs the HCA to mark the last packet currently in its receive queue, and to inform the CPU when this last packet has been written to system controller 24. At this point, the CPU can send its dummy read command and begin processing the data in the memory.

Once it is assured that all of the relevant data have reached their destination in memory 22, CPU 21 reads the cause of the current interrupt from register 56, at a cause reading step 78. Based on this information, the CPU processes the data that peripheral device 25 has placed in the memory, at a data processing step 80. Unlike methods of interrupt processing known in the art, all of the steps in the method of Fig. 3 are carried out locally, typically over busses 50 and 52, without the need for messages to traverse fabric 26.

It will be appreciated that the preferred embodiments described above are cited by way of example, and that the present invention is not limited to what has been particularly shown

and described hereinabove. Rather, the scope of the present invention includes both combinations and subcombinations of the various features described hereinabove, as well as variations and modifications thereof which would occur to persons skilled in the art upon reading the foregoing description and which are not disclosed in the prior art.

CLAIMS

1. A method for conveying data over a packet-switching network, comprising:

receiving data from a peripheral device for transmission via the network to a memory associated with a central processing unit (CPU);

receiving an interrupt signal from the peripheral device associated with the data;

sending one or more data packets containing the data over the network to a host network interface serving the memory and the CPU; and

sending an interrupt packet over the network to the host network interface, responsive to which an interrupt input of the CPU is asserted only after the one or more data packets have arrived at the host network interface.

2. A method according to claim 1, wherein receiving the data comprises receiving parallel data over a local bus from the peripheral device.

3. A method according to claim 1, wherein receiving the data comprises receiving data to be written to the memory by direct memory access.

4. A method according to claim 1, wherein sending the interrupt packet comprises reading a cause of the interrupt from the peripheral device, and incorporating the cause in the interrupt packet.

5. A method according to claim 4, and comprising receiving the interrupt packet at the host network interface, and writing the cause to a predetermined address in the memory, to be read by the CPU after the interrupt input is asserted.

6. A method according to any of the preceding claims, wherein sending the interrupt packet comprises sending the interrupt packet after receiving an acknowledgment from the memory that the data have been written thereto.

7. A method according to any of claims 1-5, wherein sending the one or more data packets comprises sending the data packets over a selected channel through the network, and wherein sending the interrupt packet comprises sending the interrupt packet over the selected channel following the data packets.

8. A method according to any of claims 1-5, and comprising:

receiving the data packets and the interrupt packet at the host network interface;

conveying the data in the packets for delivery to the memory over a local bus coupling the host network interface to the memory and the CPU; and

notifying the CPU when all of the data have been conveyed.

9. A method according to claim 8, wherein conveying the data in the packets comprises passing the data to a system controller on the bus, and wherein notifying the CPU comprises informing the CPU when an acknowledgment is received by the host network interface from the system controller.

10. A method according to claim 9, wherein informing the CPU comprises asserting the interrupt input of the CPU after the acknowledgment from the system controller has been received.

11. A method according to claim 8, wherein notifying the CPU comprises asserting the interrupt input of the CPU responsive to receiving the interrupt packet at the host network interface.

12. Network interface apparatus, comprising:

a target channel adapter, which is operative to receive data from a peripheral device for transmission via a packet-switching network to a memory associated with a central processing unit (CPU) and to send one or more data packets containing the data over the network to a host network interface serving the memory and the CPU; and

a target interface processor, adapted to receive an interrupt signal from the peripheral device associated with the data, and to send an interrupt packet over the network to the host network interface, responsive to which an interrupt input of the CPU is asserted only after the one or more data packets have arrived at the host network interface.

13. Apparatus according to claim 12, wherein the target channel adapter comprises an interface to a local parallel bus linked to the peripheral device, over which the device sends the data.

14. Apparatus according to claim 12, wherein the target channel adapter is operative to read a cause of the interrupt from the peripheral device, and wherein the processor is adapted to incorporate the cause in the interrupt packet.

15. Apparatus according to claim 14, and comprising a host channel adapter, coupled to receive the interrupt packet at the host network interface, and to write the cause to a predetermined address in the memory, to be read by the CPU after the interrupt input is asserted.

16. Apparatus according to any of claims 12-15, wherein the processor is adapted to send the interrupt packet after receiving an acknowledgment from the memory that the data have been written thereto.

17. Apparatus according to any of claims 12-15, wherein the target channel adapter is coupled to send the data packets over a selected channel through the network, and wherein the processor is adapted to send the interrupt packet over the selected channel following the data packets.

18. Apparatus according to claim 17, and comprising a switch coupling the target channel adapter and the processor to the network, wherein the switch comprises a receive queue into which the target channel adapter places the data packets, and wherein the processor is adapted to place the interrupt packet into the receive queue following the data packets.

19. Apparatus according to any of claims 12-15, and comprising a host interface unit, which is coupled to receive the data and interrupt packets transmitted over the network, and is operative to convey the data in the packets for delivery to the memory over a local bus coupled to the memory and the CPU and to notify the CPU when all of the data have been conveyed.

20. Apparatus according to claim 19, wherein the host interface unit is coupled to assert the interrupt to the CPU responsive to the interrupt packet.

21. Apparatus according to any of claims 12-15, wherein the target channel adapter comprises an InfiniBand adapter.

22. Network interface apparatus, comprising:

a host channel adapter, which is operative to receive data packets transmitted over a packet-switching network from a peripheral device, and to convey data from the packets for delivery to a memory associated with a CPU over a local bus that is coupled to the memory and the CPU, and further to receive an interrupt packet sent over the network responsive to an interrupt signal asserted by the peripheral device after sending the data to the network; and

(12) INTERNATIONAL APPLIC[redacted]N PUBLISHED UNDER THE PATENT CO[redacted]ATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
15 March 2001 (15.03.2001)

PCT

(10) International Publication Number
WO 01/18654 A1

(51) International Patent Classification[7]: G06F 13/00

(21) International Application Number: PCT/IL00/00540

(22) International Filing Date:
7 September 2000 (07.09.2000)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:

60/152,849	8 September 1999 (08.09.1999)	US
60/175,339	10 January 2000 (10.01.2000)	US
09/559,352	27 April 2000 (27.04.2000)	US

(71) Applicant *(for all designated States except US)*: **MELLANOX TECHNOLOGIES LTD.** [IL/IL]; P.O. Box 86, 20692 Yokneam (IL).

(72) Inventors; and

(75) Inventors/Applicants *(for US only)*: **KAGAN, Michael** [IL/IL]; Hashomer Street 71, 30900 Zichron Yaakov (IL). **CRUPNICOFF, Diego** [AR/IL]; Sitvanit Street 9, 34793 Haifa (IL). **GABBAY, Freddy** [IL/IL]; Derech Hashalom Street 75/2, 67942 Tel Aviv (IL). **ROTTENBERG, Shimon** [IL/IL]; Dubnov Street 27, 64597 Tel Aviv (IL).

(74) Agents: **COLB, Sanford, T.** et al.; Sanford T. Colb & CO., P.O. Box 2273, 76122 Rehovot (IL).

(81) Designated States *(national)*: AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

[Continued on next page]

(54) Title: SYNCHRONIZATION OF INTERRUPTS WITH DATA PACKETS

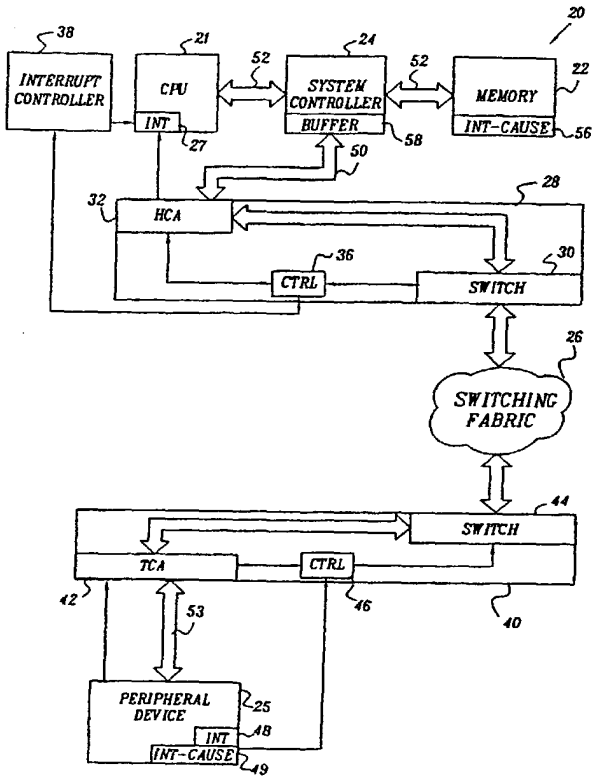

(57) Abstract: A method and apparatus for conveying data over a packet-switching network (26). Data are received from a peripheral device (25) for transmission via the network to a memory (22) associated with a central processing unit (CPU) (21), followed by an interrupt signal from the peripheral device associated with the data. One or more data packets containing the data are sent over the network to a host network interface (32) serving the memory and the CPU, followed by an interrupt packet sent over the network to the host network interface. Responsive to the interrupt packet, an interrupt input of the CPU is asserted only after the one or more data packets have arrived at the host network interface.

PCT appln. no. filed (published)

Express Mail Number
EV 009953392 US

WO 01/18654 A1

(84) **Designated States** *(regional)*: ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published:
— *With international search report.*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

a host interface processor, adapted, responsive to the interrupt packet, to notify the CPU when all of the data have been conveyed to the local bus.

23. Apparatus according to claim 22, wherein the host channel adapter is operative to convey the data to the memory by direct memory access.

24. Apparatus according to claim 22, wherein the host channel adapter is operative to convey the data to a system controller on the bus, and wherein the CPU is notified when an acknowledgment is received by the host channel adapter from the system controller.

25. Apparatus according to claim 24, wherein the host interface processor is coupled to assert the interrupt input of the CPU after the acknowledgment from the system controller has been received.

26. Apparatus according to any of claims 22-25, wherein the host interface processor is coupled to assert the interrupt input of the CPU responsive to receipt of the interrupt packet at the host network interface.

27. Apparatus according to any of claims 22-25, wherein the host channel adapter comprises an InfiniBand adapter.

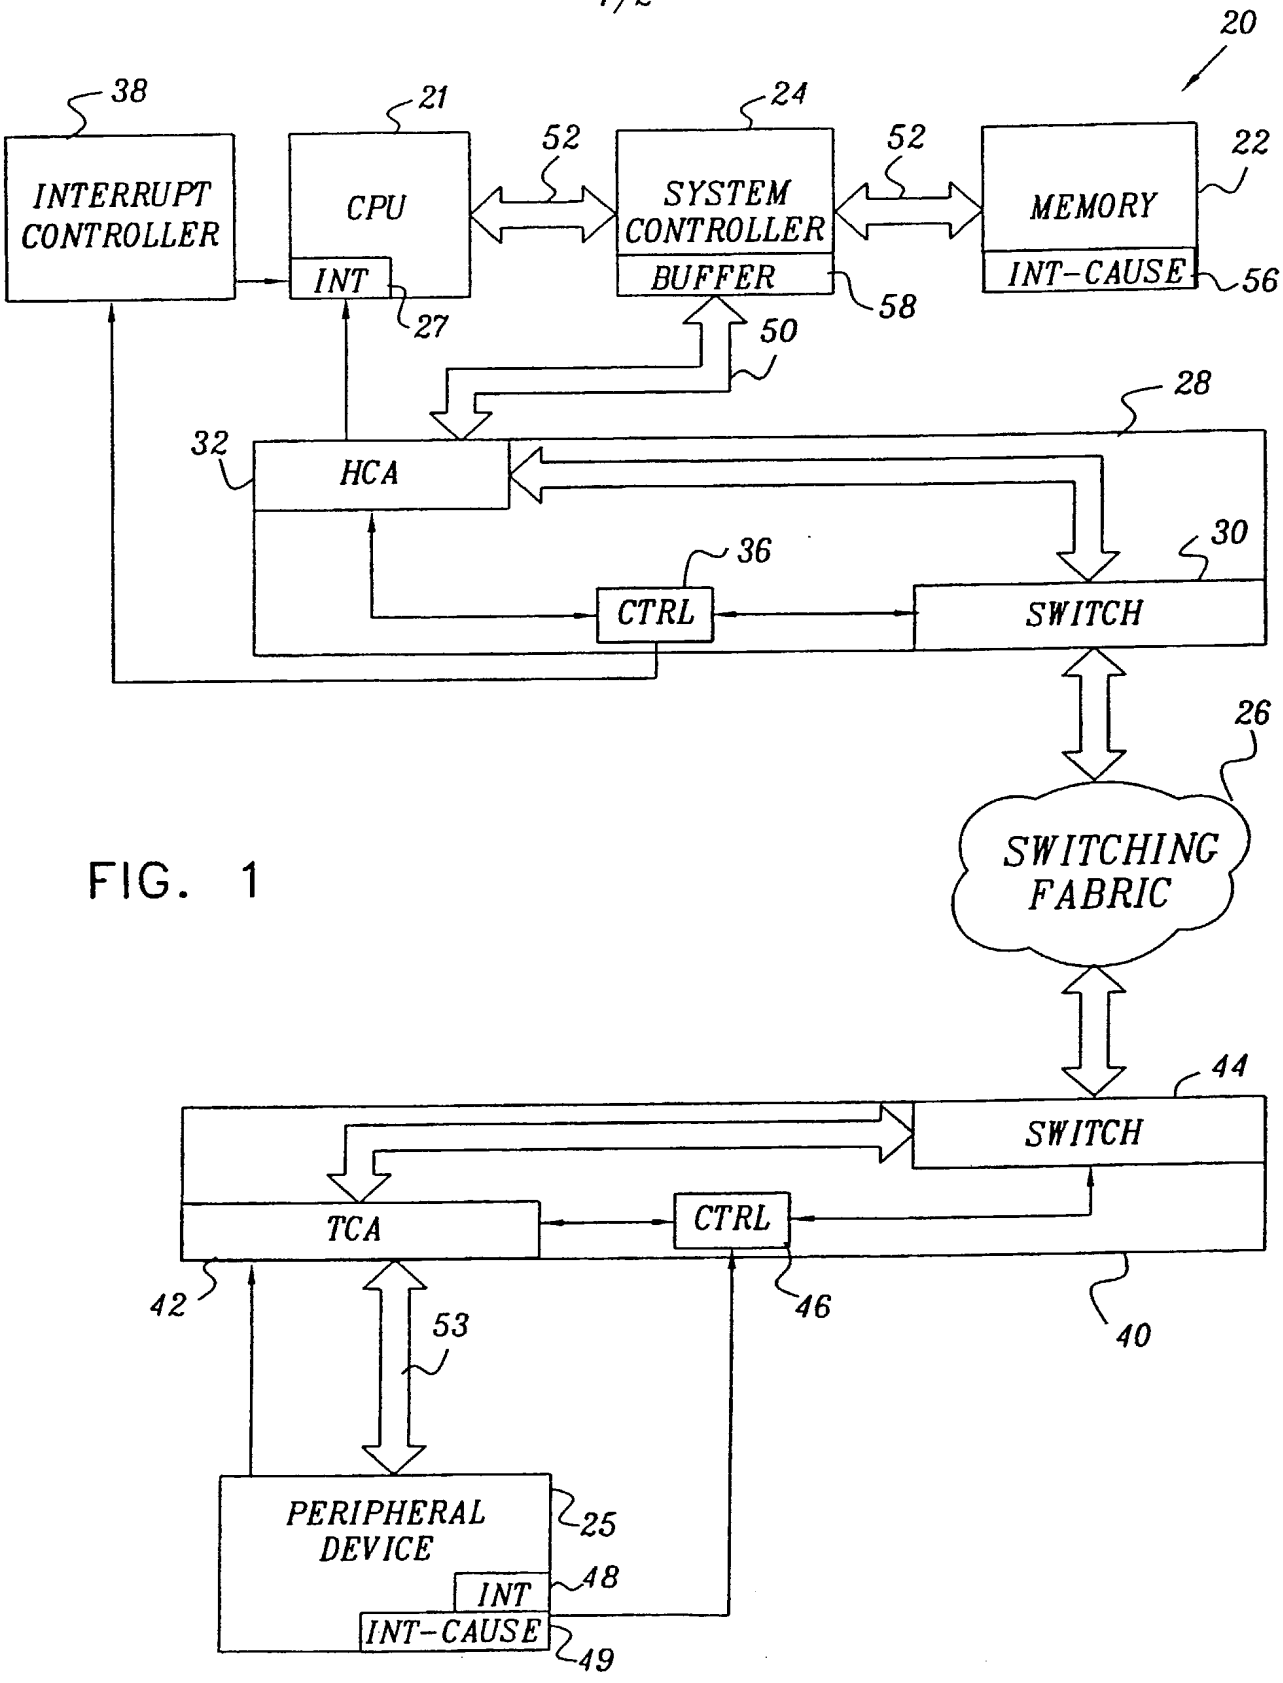
20
38
INTERRUPT CONTROLLER
21
CPU
INT
27
52
24
SYSTEM CONTROLLER
BUFFER
58
52
22
MEMORY
INT-CAUSE
56
50
28
32
HCA
36
CTRL
30
SWITCH
26
SWITCHING FABRIC
44
SWITCH
TCA
CTRL
42
46
40
53
PERIPHERAL DEVICE
25
INT
48
INT-CAUSE
49

FIG. 1

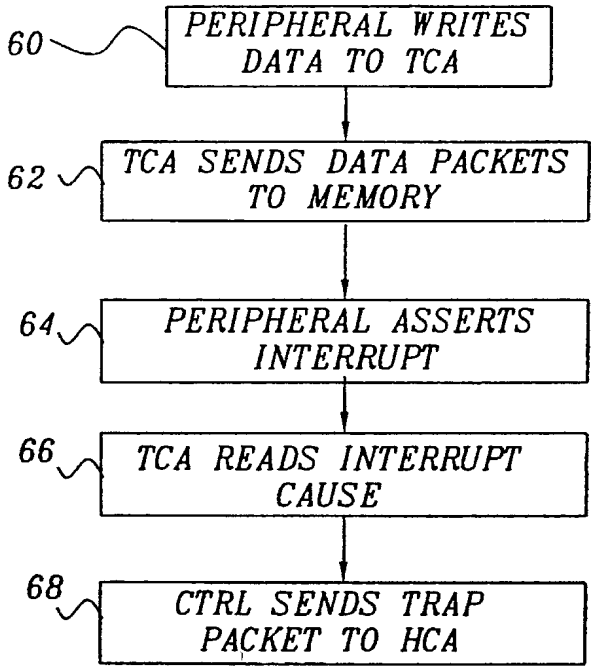
60
PERIPHERAL WRITES DATA TO TCA
62
TCA SENDS DATA PACKETS TO MEMORY
64
PERIPHERAL ASSERTS INTERRUPT
66
TCA READS INTERRUPT CAUSE
68
CTRL SENDS TRAP PACKET TO HCA

FIG. 2

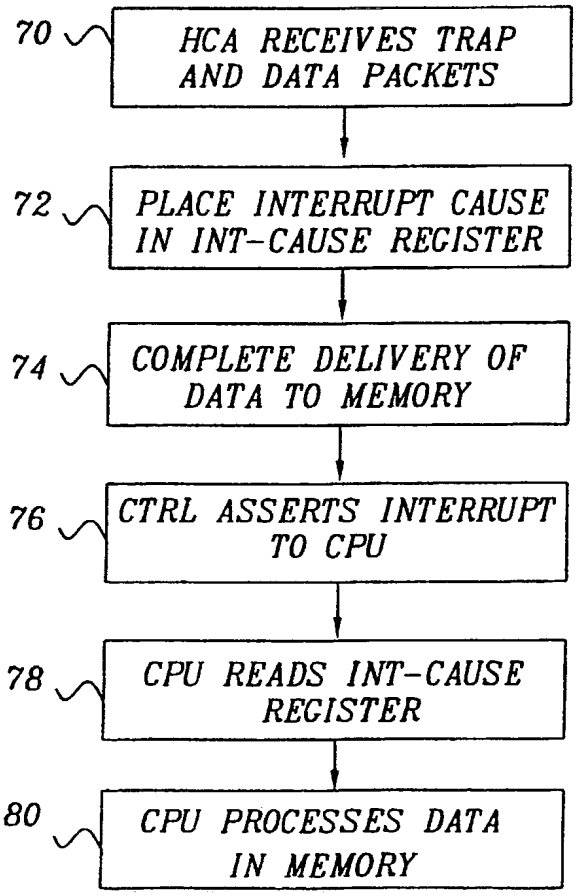
70
HCA RECEIVES TRAP AND DATA PACKETS
72
PLACE INTERRUPT CAUSE IN INT-CAUSE REGISTER
74
COMPLETE DELIVERY OF DATA TO MEMORY
76
CTRL ASSERTS INTERRUPT TO CPU
78
CPU READS INT-CAUSE REGISTER
80
CPU PROCESSES DATA IN MEMORY

FIG. 3

Attorney
Docket No.: COLB-121XX

DECLARATION AND POWER OF ATTORNEY

As a below-named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name;

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

SYNCHRONIZATION OF INTERRUPTS WITH DATA PACKETS

The specification of which (check one):

[ ] **is attached hereto.** [X] **was filed on** March 8, 2002 **as Application No.** 10/070,594 ; **amended on** _____ **(if applicable).**

[X] was filed as PCT International. Appl. No. PCT/IL00/00540 on 07 September 2000, and was amended under PCT Article 19 on _____ (if applicable).

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with Title 37, Code of Federal Regulations §1.56(a).

I hereby claim foreign priority benefits under Title 35, USC §119(a)-(d) of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s)		Date Filed	Priority Claimed	
(Number)	(Country)	(Day/Month/Year)	[ ] Yes	[ ] No
(Number)	(Country)	(Day/Month/Year)	[ ] Yes	[ ] No
(Number)	(Country)	(Day/Month/Year)	[ ] Yes	[ ] No

I hereby claim the benefit under Title 35, USC §119(e) of any United States provisional application(s) listed below:

60/152,849 (Application Number)	08 September 1999 (Filing Date)
60/175,339 (Application Number)	10 January 2000 (Filing Date)
(Application Number)	(Filing Date)

Express Mail Number

EV 009950246 US

45054

Sheet 2 of 5

Attorney
Docket No.: COLB-121XX

I hereby claim the benefit under Title 35 USC §120 of any United States application(s) listed below and insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35 USC §112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, §1.56(a) which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

(Application No.)	(Filing Date)	(Patented/pending/abandoned)
09/559,352	27 April 2000	

POWER OF ATTORNEY: As a named inventor, I hereby appoint the following attorney(s) to prosecute this application and transact all business connected therewith in the Patent and Trademark Office, and to file with the USRO any International Application based thereon.

(8)

Stanley M. Schurgin, Reg. No. 20,979
Charles L. Gagnebin III, Reg. No. 25,467
Victor B. Lebovici, Reg. No. 30,864
[illegible]verly E. Hjorth, Reg. No. 32,033
Holliday C. Heine, Reg. No. 34,346
Gordon R. Moriarty, Reg. No. 38,973
James F. Thompson, Reg. No. 36,699
David A. Dagg, Reg. No. 37,809

Address all correspondence to:
Customer Number: 207

WEINGARTEN, SCHURGIN, GAGNEBIN & LEBOVICI LLP
Ten Post Office Square
Boston, Massachusetts 02109
Telephone: (617) 542-2290
Telecopier: (617) 451-0313

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

1-00

Full Name of First/Sole Inventor: Michael Kagan		
City of Residence Zichron Yaakov	State or Country Israel ILK	Country of Citizenship Israel
Post Office Address Hashomer Street 71	City 30900 Zichron Yaakov	State or Country Zip Code Israel
Signature: (Please sign and date in permanent ink.) X		Date signed: X 22-05-02

4505

Sheet 3 of 5

Attorney
Docket No.: COLB-121XX

2-00

Full Name of Second/Joint Inventor: Diego Crupnicoff		
City of Residence Haifa	State or Country Israel IL	Country of Citizenship Argentina
Post Office Address Sitvanit Street 9	City 34793 Haifa	State or Country Zip Code Israel
Signature: (Please sign and date in permanent ink.) X		Date signed: X MAY 19, 2002

01

Attorney
Docket No.: COLB-121XX

3-00

Full Name of Third/Joint Inventor: Freddy Gabbay		
City of Residence Tel Aviv	State or Country Israel ILL	Country of Citizenship Israel
Post Office Address Derech Hashalom Street 75/2	City 67942 Tel Aviv	State or Country Zip Code Israel
Signature: (Please sign and date in permanent ink.) X F. Gab		**Date signed:** X 21/5/2002

400

Full Name of Fourth/Joint Inventor: Shimon Rottenberg		
City of Residence Tel Aviv	State or Country Israel ILX	Country of Citizenship Israel
Post Office Address Dubnov Street 27	City 64597 Tel Aviv	State or Country Zip Code Israel
Signature: (Please sign and date in permanent ink.) X [signature]		**Date signed:** X MAY 21 2002

CLG/mc/268666-1